

Assistive Planning in Complex, Dynamic Environments: a Probabilistic Approach

Pete Trautman
 Galois Inc., Portland OR
 peter.trautman@galois.com

Abstract—We explore the probabilistic foundations of shared control in complex dynamic environments. In order to do this, we formulate shared control as a random process and describe the joint distribution that governs its behavior. For tractability, we model the relationships between the operator, autonomy, and crowd as an undirected graphical model. Further, we introduce an interaction function between the operator and the robot, that we call “agreeability”; in combination with the methods developed in [23], we extend a cooperative collision avoidance autonomy to shared control. We therefore quantify the notion of simultaneously optimizing over agreeability (between the operator and autonomy), and safety and efficiency in crowded environments. We show that for a particular form of interaction function between the autonomy and the operator, linear blending is recovered exactly. Additionally, to recover linear blending, unimodal restrictions must be placed on the models describing the operator and the autonomy. In turn, these restrictions raise questions about the flexibility and applicability of the linear blending framework. Additionally, we present an extension of linear blending called “operator biased linear trajectory blending” (which formalizes some recent approaches in linear blending such as [9]) and show that not only is this also a restrictive special case of our probabilistic approach, but more importantly, is statistically unsound, and thus, mathematically, unsuitable for implementation. Instead, we suggest a statistically principled approach that guarantees data is used in a consistent manner, and show how this alternative approach converges to the full probabilistic framework. We conclude by proving that, in general, linear blending is suboptimal with respect to the joint metric of agreeability, safety, and efficiency.

I. INTRODUCTION

Fusing human and machine capabilities has been an active research topic in computer science for decades. In robotics related applications, this line of research is often referred to as *shared autonomy*. While shared autonomy has been addressed *in toto*, it can be broken down into more specialized areas of research. In particular, some researchers focus on the fusing of human and machine “perception” (see [18]). Similarly, a great deal of research has focused on fusing human and machine “decision making” in the machine learning [29], [16], control theory [11], and human robot interaction communities [13].

Under the umbrella of shared decision making, an even more focussed line of research has emerged: shared control, whereby the moment to moment control commands sent to the platform motors are a synthesis of human input and autonomy input. Broadly speaking, shared control has been deployed in two cases: shared teleoperation (where the human operator is not co-located with the robot) and onboard shared control (where the human operator is physically on the robot). Shared teleoperation is used for numerous applications: search and rescue [4] and extraterrestrial robotics [2] are two examples.

Onboard shared control has likewise enjoyed a wide variety

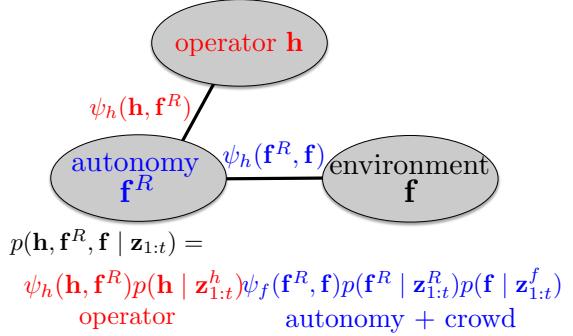


Fig. 1: Diagram depicting the relationships of probabilistic shared control.

of use cases: assistive wheelchair technology [12], assistive automobile driving, and assistive manufacturing vehicle operation (e.g., forklifts) are just a few of the examples. For any of these cases, shared control can be broken down into an autonomous modeling step, a human modeling step, and a human-machine arbitration step (see [9] for a compelling argument justifying this classification scheme).

In this paper, we explore the probabilistic foundations of onboard shared control in the presence of dynamic obstacles (“the crowd”): Figure 1 depicts our approach. In order to do this, we formulate shared control as a random process and describe the joint distribution (over the random operator, autonomy, and crowd functions) that governs behavior and then propose a tractable model captures notions of operator-autonomy agreeability, safety, and efficiency; further, we prove that linear blending is a special case of this approach. In the course of the proof, we observe the following: first, in linear blending, the autonomy reasons *independently* of the operator during optimization, so even if we had a precise operator model, the autonomy would not be informed of this information until after the optimization. Second, in linear blending, the autonomy is limited to a single optimal decision, which is then averaged with the operator input—this approach leaves no flexibility in how the autonomy might assist the user. Next, we present an extension of linear blending that can accommodate more than one statistic about the user (e.g., user inferred goal information, preferred trajectories, etc), and find that state of the art approaches to this problem are prone to statistical inconsistency. We thus present a statistically valid model for how to properly condition the autonomy on user statistics. We conclude with a section on the optimality (with respect to agreeability, safety, and efficiency) of linear blending and our probabilistic approach.

II. RELATED WORK

We begin by defining the arbitration step of linear blending:

$$u_{LB}^s(t) = K_h u_t^h + K_R u_{t+1}^R, \quad (\text{II.1})$$

where, at time t , $u_{LB}^s(t)$ is the linearly blended shared control command sent to the platform actuators, u_t^h is the human operator input (joystick deflections, keyboard inputs, etc.), u_{t+1}^R is the next autonomy command, and K_h, K_R are the operator and autonomy *arbitration functions*, respectively. To ensure that the magnitude of $u_{LB}^s(t)$ does not exceed the magnitude of u_t^h or u_{t+1}^R , we require that $K_h + K_R = 1$. We observe that since $K_h = 1 - K_R$, there is only one free parameter in this formulation.

This linear arbitration model has enjoyed wide adoption in the assistive wheelchair community ([5], [27], [17], [26], [30], [19], [25], [15]). Outside of the wheelchair community, shared control path planning researchers have widely adopted Equation II.1 as a *de-facto* standard protocol, as extensively argued in [9], [8] (in [9], it is argued that “linear policy blending can act as a common lens across a wide range of literature”). Additionally, the work of [20], [28] advocates the broad adoption of a linear arbitration step for shared control.

For the purposes of this article, we describe how each quantity of Equation II.1 is computed:

1) Collect the data at time t : $u_{1:t}^h$ are the historical operator inputs. $\mathbf{z}_{1:t}^R$ are the historical measurements of the state of the robot (such as odometry, localization, SLAM output, etc).

2) Compute the *autonomous* input u_{t+1}^R : This quantity may be computed using nearly any off the shelf planning algorithm, and is dependent on application. The “Dynamic Window Approach” [10] and “Vector Field Histograms+” [24] are popular approaches to perform obstacle avoidance for wheelchairs. Sometimes, the autonomy is biased according to data about the operator—for instance, one might imagine an offline training phase where the robot is taught “how” to move through the space, and then this data could be agglomerated using, e.g., inverse optimal control. Alternatively, one might bias the autonomous decision making by conditioning the planner on the predicted or known human goal.

3) Compute the arbitration parameters K_h and K_R . A wide variety of heuristics have been adopted to compute this parameter: confidence in robot trajectory, smoothness, mitigating jerk, operator reliability, user desired trajectories, safeguarding against unsafe trajectories, etc. Indeed, much of the shared control literature is devoted to developing novel heuristics to compute this parameter.

4) Compute the shared control $u_{LB}^s(t)$ using Equation II.1.

Typically, the data $u_{1:t}^h$ are interpreted literally—no likelihood or predictive model filters this data stream. In other approaches, operator intention is modeled using a combination of dynamic Bayesian networks or Gaussian mixture models.

For this paper, we adopt the notation $\mathbf{z}_t^h \doteq u_t^h$, (that is, we treat operator inputs as *measurements* of the operator *trajectory*, $\mathbf{h}: t \in \mathbb{R} \rightarrow \mathcal{X}$, where \mathcal{X} is the action space). Similarly, we define measurements $\mathbf{z}_{1:t}^R$ of the robot trajectory \mathbf{f}^R and measurements $\mathbf{z}_{1:t}^i$ of the i 'th static or dynamic obstacle

trajectory \mathbf{f}^i . We thus work in the space of distributions over the operator function \mathbf{h} , autonomy function \mathbf{f}^R , and crowd function $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^{n_t})$, measured through $\mathbf{z}_{1:t}^f$. The integer n_t is the number of people in the crowd at time t .

We comment on the work of [6], [14], [7]. In these papers, the authors construct a probabilistic model over user forward trajectories (i.e., a specific and personalized instantiation of $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$). They then formulate shared control as a *partially observable Markov decision process* (POMDP), in order to capture the effects of robot actions on the probabilistic model of the operator. However, for tractability of the POMDP, the autonomy is limited to only choosing from the next available state—this makes assistive navigation through crowds impossible (as shown in [21]). Further, the potential autonomous actions are limited by the input device; in one application, the autonomy was only able to reason over 9 directions. As we discuss in Section VI, limiting the autonomy to such a degree can negatively impact performance.

III. FOUNDATIONS AND IMPLEMENTATION OF PROBABILISTIC SHARED CONTROL

As stated in Section I, we seek to explore the probabilistic foundations of assistive onboard shared control; we thus posit

$$\begin{aligned} u_{PSC}^s(t) &= \mathbf{f}_{t+1}^{R*} \\ (\mathbf{h}, \mathbf{f}^R, \mathbf{f})^* &= \arg \max_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}} p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} | \mathbf{z}_{1:t}^h, \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f). \end{aligned} \quad (\text{III.1})$$

That is, probabilistic shared control (PSC) is the *maximum a-posteriori* (MAP) value of the joint distribution over the operator, autonomy, and crowd. We suggest this approach partly based on what we have learned from fully autonomous navigation in human crowds [23]—that the *interaction* model between the robot and the human is the most important quantity—and partly based on the following: by formulating shared control as the MAP value of a joint probability distribution, we can explore the modeling limitations imposed by linear blending and the performance consequences of these limitations.

In the remainder of this section, we explain how the approach in Equation III.1 is a natural extension of our previous work in [22], and present a tractable model of the joint distribution over the operator, autonomy, and crowd. The probabilistic graphical model in Figure 2 guides our derivation.

Definition 1: A cooperative human crowd navigation model (as in [22]) is described by

$$p(\mathbf{f}^R, \mathbf{f} | \bar{\mathbf{z}}_{1:t}) = \psi(\mathbf{f}^R, \mathbf{f}) p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) \prod_{i=1}^{n_t} p(\mathbf{f}^i | \mathbf{z}_{1:t}^i) \quad (\text{III.2})$$

where $\bar{\mathbf{z}}_{1:t} = [\mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f]$ and $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R), p(\mathbf{f}^i | \mathbf{z}_{1:t}^i)$ are the robot and crowd individual dynamical prediction functions, respectively. The robot-crowd interaction function $\psi(\mathbf{f}^R, \mathbf{f})$ captures joint collision avoidance: that is, how the robot and the crowd cooperatively move around one another so that the other can pass. The graphical model is presented on the left hand side of Figure 2.

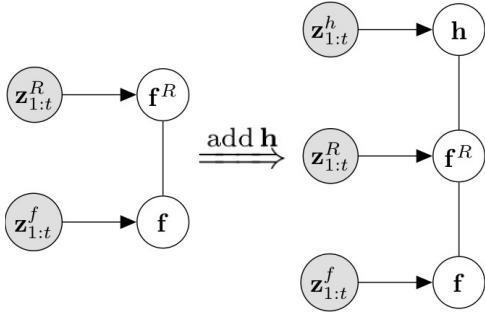


Fig. 2: Correspondence between the cooperative model of Equation III.2 (left hand side of figure) and probabilistic shared control in Equation III.3.

Definition 2: A probabilistic shared control (PSC) model in the presence of static or dynamic obstacles is

$$p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) = \psi(\mathbf{h}, \mathbf{f}^R)p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) \quad (\text{III.3})$$

where $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ is the predictive distribution over the operator, $\mathbf{z}_{1:t}^h$ is data generated by the operator, and $\psi(\mathbf{h}, \mathbf{f}^R)$ is the interaction function between the operator and the robot. In analogy with [22], the interaction function $\psi(\mathbf{h}, \mathbf{f}^R)$ intends to capture how “agreeable” the robot trajectory is with the desires of the operator. In this article, we choose $\psi(\mathbf{h}, \mathbf{f}^R) = \exp(-\frac{1}{2\gamma}(\mathbf{h} - \mathbf{f}^R)(\mathbf{h} - \mathbf{f}^R)^\top)$, where γ captures how tightly the autonomy must couple to the operator; it is important to point out that many other reasonable functions could be used here. Importantly, as shown in Theorem 4.1, this is the (implicit) operator-robot interaction function of linear blending. The graphical model is presented on the right hand side of Figure 2.

The model in Equation III.3 can be understood in the following way: the autonomy (modeled with $p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$) optimizes over safety and efficiency in crowds—specifically, the robot-crowd interaction function $\psi(\mathbf{f}^R, \mathbf{f})$ models how to move *through* a crowd in the safest and most efficient way. Similarly, the function $\psi(\mathbf{h}, \mathbf{f}^R)$ models how “agreeable” the robot path \mathbf{f}^R is to the operator path \mathbf{h} . By formulating \mathbf{h} and \mathbf{f}^R as random functions, flexibility of both parties is modeled, thus enabling a more realistic arbitration (where both parties are willing to compromise). This idea was captured in [21]: the robot and the crowd were treated as random functions so that they could come to a compromise about how to (cooperatively) pass by one another.

Thus, by choosing the joint $\arg \max_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}}$ as in Equation III.1, we *simultaneously* optimize over three quantities: agreeability, safety, and efficiency. We suggest that such a joint optimization over safety and efficiency (the domain of the autonomy) and agreeability (adhering to the desires of the operator) is a reasonable way to formulate shared control. Our analysis in the following sections provides credibility to this presumption.

IV. LINEAR BLENDING AS PROBABILISTIC SHARED CONTROL

To understand linear blending as a special case of Equation III.1, we consider the following conditioning relationship:

$$p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) = p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f, \mathbf{h})p(\mathbf{h} \mid \mathbf{z}_{1:t}^h). \quad (\text{IV.1})$$

To understand how the linear blending modeling assumptions effect the full joint, we first insert the linear blending *operator* model: $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h) = \delta(\mathbf{h} - \mathbf{z}_t^h)$. Thus,

$$\begin{aligned} p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) &= p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \mathbf{h})\delta(\mathbf{h} - \mathbf{z}_t^h) \\ &= p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \mathbf{z}_t^h). \end{aligned}$$

In this case,

$$\arg \max_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}} p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) = \arg \max_{\mathbf{f}^R, \mathbf{f}} p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \mathbf{z}_t^h).$$

Because this distribution is already conditioned on \mathbf{z}_t^h , there is no need for a linear arbitration step (we prove this in Theorem 4.1). Further, by fixing the operator at \mathbf{z}_t^h , we are no longer *jointly* optimizing over the autonomy and the operator.

Theorem 4.1 (Equation III.1 generalizes linear blending): Let $p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \mathbf{z}_t^h) = \psi_h(\mathbf{z}_t^h, \mathbf{f}_{t+1}^R)p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$ and use Laplace’s Approximation [3] to approximate $p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$, where \mathbf{f}_{t+1}^R is a mode of the distribution:

$$p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) = \mathcal{N}(\mathbf{f}_{t+1}^R \mid \bar{\mathbf{f}}_{t+1}^R, \boldsymbol{\sigma}^R).$$

Then $u_{PSC}^s(t) = \boldsymbol{\sigma}(\frac{1}{\gamma}\mathbf{z}_t^h + \frac{1}{\sigma^R}\bar{\mathbf{f}}^R)$, $\boldsymbol{\sigma}^{-1} = (\gamma^{-1} + (\boldsymbol{\sigma}^R)^{-1})$, and γ is the operator-autonomy attraction parameter.

Proof: We compute

$$\begin{aligned} u_{PSC}^s(t) &= \arg \max_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}} p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) \\ &\propto \arg \max_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}} \mathcal{N}(\mathbf{f}_{t+1}^R \mid \mathbf{z}_t^h, \gamma)\mathcal{N}(\mathbf{f}_{t+1}^R \mid \bar{\mathbf{f}}_{t+1}^R, \boldsymbol{\sigma}^R) \\ &= \arg \max_{\mathbf{f}_{t+1}^R} \mathcal{N}(\mathbf{f}_{t+1}^R \mid \boldsymbol{\mu}, \boldsymbol{\sigma}) \\ &= \boldsymbol{\sigma}\left(\frac{1}{\gamma}\mathbf{z}_t^h + \frac{1}{\sigma^R}\bar{\mathbf{f}}^R\right) \end{aligned} \quad (\text{IV.2})$$

where $\boldsymbol{\mu} = \boldsymbol{\sigma}\left(\frac{1}{\gamma}\mathbf{z}_t^h + \frac{1}{\sigma^R}\bar{\mathbf{f}}^R\right)$ and $\boldsymbol{\sigma}^{-1} = (\gamma^{-1} + (\boldsymbol{\sigma}^R)^{-1})$. ■

To make the correspondence with linear blending explicit, we choose $\boldsymbol{\sigma}^R = \gamma/K_R - \gamma$ and $\gamma = \boldsymbol{\sigma}^R/K_h - \boldsymbol{\sigma}^R$. Then

$$\begin{aligned} \boldsymbol{\sigma}\left(\frac{1}{\gamma}\mathbf{z}_t^h + \frac{1}{\sigma^R}\bar{\mathbf{f}}^R\right) &= \frac{\boldsymbol{\sigma}^R}{\boldsymbol{\sigma}^R + \gamma}\mathbf{z}_t^h + \frac{\gamma}{\boldsymbol{\sigma}^R + \gamma}\bar{\mathbf{f}}^R \\ &= K_h\mathbf{z}_t^h + K_R\bar{\mathbf{f}}^R. \end{aligned}$$

We recall that $K_h + K_R = 1$ and so the mixing of $\boldsymbol{\sigma}_R$ and γ is expected. Only one free parameter exists in linear blending.

We now examine the how the assumptions of linear blending can effect performance. The linear blending restriction of the robot-crowd distribution to a unimodal Gaussian during arbitration $p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f) = \mathcal{N}(\mathbf{f}_{t+1}^R \mid \bar{\mathbf{f}}_{t+1}^R, \boldsymbol{\sigma}_{t+1}^R)$ can lead to severe restrictions in shared control capability; even though a locally optimal autonomy strategy \mathbf{f}^R is included in the linear blend $u_{LB}^s = K_h\mathbf{z}_t^h + K_R\bar{\mathbf{f}}^R$, there is no guarantee that the human operator will choose this optima, or even a nearby optimum (see Section VI). Alternatively, by maintaining the multitude of hypotheses inherent to $p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$, we greatly increase the possibility that the autonomy will be able to assist the operator in a way that is both desirable and safe. As an example, linear trajectory blending over a *safe* operator input and a *safe* autonomous input can result in an unsafe shared trajectory. That is, the weighted average of two safe trajectories can be averaged into an unsafe trajectory.

V. CONDITIONAL TRAJECTORY BLENDING

In this section, we extend our definition of linear blending to capture the salient characteristics of recent approaches such as [9]. Broadly speaking, this line of work seeks to include additional data about the operator *before* arbitration. We show that this is statistically unfounded, and thus introduce a method (conditional trajectory blending) that incorporates operator information in a statistically principled way.

Definition 3 (Extending Equation II.1): Let

- 1) $\bar{\mathbf{h}} \sim p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ be defined through time $T > t$, such that $\bar{\mathbf{h}}: [1, T] \subset \mathbb{R} \mapsto \mathcal{X}$,
- 2) $\mathbf{f}^R \sim p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f)$, with $\bar{\mathbf{f}}^R: [1, T] \subset \mathbb{R} \mapsto \mathcal{X}$,
- 3) u_{LTB}^s be the shared control $u_{LTB}^s: [1, T] \subset \mathbb{R} \mapsto \mathcal{X}$,
- 4) and K_h, K_R be the arbitration functions.

Definition 4 (LTB): Linear Trajectory Blending (LTB) is defined as 1) Sample $\bar{\mathbf{f}}^R = \arg \max_{\mathbf{f}^R, \mathbf{f}} p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f)$, 2) Sample $\bar{\mathbf{h}} = \arg \max_{\mathbf{h}} p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$, and 3) Construct the shared control $u_{LTB}^s = K_h \mathbf{h} + K_R \mathbf{f}^R$.

By extending to trajectories, we can more easily incorporate operator information into the arbitration step. In this vein, we now define an extension of linear trajectory blending that biases the autonomous decision making on operator data (motivated by the approach in [9]).

Definition 5 (LTBo): Let $p(G \mid \mathbf{z}^h)$ be a distribution about the operator, where $\mathbf{z}^h \subseteq \mathbf{z}_{1:t}^h$. Then *operator biased linear trajectory blending* (LTBo) is defined as

- 1) Sample $G_1 \sim p(G \mid \mathbf{z}^h)$. G_1 could be a trajectory, a waypoint, a goal, or any other relevant quantity.
- 2) Sample the *operator biased* distribution

$$\bar{\mathbf{f}}^R_{\mathbf{h}} = \arg \max_{\mathbf{f}^R, \mathbf{f}} p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f, G_1).$$

- 3) Sample the operator trajectory $\bar{\mathbf{h}} = \arg \max_{\mathbf{h}} p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$.
- 4) Construct the shared control $u_{LTBo}^s = K_h \mathbf{h} + K_R \bar{\mathbf{f}}^R_{\mathbf{h}}$.

With these definitions, we extend the approach of Equation II.1 to include a model of the operator $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ and enable seeding the autonomy with operator information. By modeling the operator with a distribution, we potentially bypass the issue of noisy inputs leading to “jittery” linear blends. By seeding the autonomy with operator statistics, we might drive u_{LTBo}^s towards solutions more closely aligned with user desire.

While this approach is *sensible*, no *principles* motivate such an extension. We thus pause to examine this approach in the context of the full joint distribution. In particular, recall Equation IV.1; if we only sample the single operator statistic $\bar{\mathbf{h}} = \arg \max_{\mathbf{h}} p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$, then we are implicitly making the assumption that $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h) = \delta(\mathbf{h} - \bar{\mathbf{h}})$; thus we have

$$p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) = p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \bar{\mathbf{h}}).$$

In this case,

$$\arg \max_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}} p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) = \arg \max_{\mathbf{f}^R, \mathbf{f}} p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \bar{\mathbf{h}}),$$

and so (as before) there is no *need* for a linear arbitration step—we have already conditioned the autonomy on the operator, and we can recover linear blending by choosing

$p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \bar{\mathbf{h}})$ as in Theorem 4.1. If we have a *separate* model $p(G \mid \mathbf{z}^h)$ about the operator, we know that it does not contain information beyond what is available in $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$, since $\mathbf{z}^h \subseteq \mathbf{z}_{1:t}^h$ (see the discussion on the “data processing inequality” in [?]); worse, $p(G \mid \mathbf{z}^h)$ can lead to statistical inconsistencies (see Lemma 5.2). We now show how to resolve this inconsistency and how to combine multiple operator data points in a statistically sound manner.

Definition 6 (Conditional Trajectory Blending):

Assume that we have the distributions $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ and $p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \mathbf{h}_b)$. Then let $\{\mathbf{h}^b\}_{b=1}^{N_h} \sim p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ be a collection of N_h samples of the operator model. If we take the model of the operator to be $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h) = \sum_{b=1}^{N_h} w^b \delta(\mathbf{h} - \mathbf{h}^b)$, then

$$\begin{aligned} p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) &= p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \mathbf{h}) \sum_{b=1}^{N_h} w^b \delta(\mathbf{h} - \mathbf{h}^b) \\ &= \sum_{b=1}^{N_h} w^b p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \mathbf{h}^b) \end{aligned} \quad (\text{V.1})$$

where $w^b = p(\mathbf{h}^b = \mathbf{h} \mid \mathbf{z}_{1:t}^h)$ is the probability of sample \mathbf{h}^b . We interpret the shared control to be

$$u_{CTB}^s = \arg \max_{\mathbf{f}^R, \mathbf{f}} \sum_{b=1}^{N_h} w^b p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \mathbf{h}^b).$$

In particular, we revisit the case of LTBo: suppose that we sample $G_1 \sim p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ — G_1 is present in this distribution since it contains all the data—and then sample $\bar{\mathbf{h}} \sim p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$. Then conditional trajectory blending tells us that we should find the $\arg \max$ of the distribution

$$\begin{aligned} u_{CTB}^s &= \arg \max_{\mathbf{f}^R, \mathbf{f}} \left[w^1 p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, G_1) + \right. \\ &\quad \left. w^2 p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \bar{\mathbf{h}}) \right]. \end{aligned}$$

In this case, then, $u_{CTB}^s \neq u_{LTBo}^s$; however, since conditional trajectory blending is derived directly from the full joint we know that it is combining the data G_1 and $\bar{\mathbf{h}}$ in a statistically sound manner. The next theorem provides information about the limiting behavior of conditional trajectory blending.

Theorem 5.1 (CTB approximates PSC): As the number of operator samples tends to infinity, probabilistic shared control (Equation III.1) is recovered.

$$\begin{aligned} \text{Proof:} \quad &\text{Representing } p(\mathbf{h} \mid \mathbf{z}_{1:t}^h) = \sum_{b=1}^{\infty} w^b \delta(\mathbf{h} - \mathbf{h}^b), \\ &p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) = p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f, \mathbf{h}) p(\mathbf{h} \mid \mathbf{z}_{1:t}^h) \\ &= p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f, \mathbf{h}) \sum_{b=1}^{\infty} w^b \delta(\mathbf{h} - \mathbf{h}^b) \\ &= \sum_{b=1}^{\infty} w^b p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f, \mathbf{h}^b). \end{aligned}$$

It is important to emphasize that conditioning the autonomy on operator statistics and then averaging $\bar{\mathbf{f}}^R_{\mathbf{h}}$ with a separate operator statistic $\bar{\mathbf{h}}$ is not just unnecessary, but potentially

statistically unsound as well, since it is unclear how such an approach originates from Equation IV.1.

Lemma 5.2 (LTBo statistically unsound): LTBo is not guaranteed to incorporate data in a statistically sound manner.

Proof: Suppose that one were to sample $\bar{\mathbf{h}} \sim p(G \mid \mathbf{z}^h)$, then sample $\mathbf{h} = \arg \max_{\mathbf{h}} p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$, then compute $\mathbf{f}^R \mathbf{h}$, and then compute $K_h \bar{\mathbf{h}} + K_R \mathbf{f}^R \mathbf{h}$. Since we have incorporated $\bar{\mathbf{h}}$ twice in the linear blend, the data has been overused. One could potentially compensate for this by “removing” the effect of double usage of $\bar{\mathbf{h}}$ in K_h , but it is unclear how to do this in a statistically sound manner.

Even if $\bar{\mathbf{h}}$ is sampled twice in CTB, it is only counted once: $u_{CTB}^s = \arg \max_{\mathbf{f}^R, \mathbf{f}} [2w^1 p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \bar{\mathbf{h}})]$. ■

Thus, if we have a model of $p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f, \mathbf{h}^b)$ and a model of the operator $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$, we have a clear mandate for how to correctly formulate shared control.

VI. OPTIMALITY OF SHARED CONTROL

We start with the following Gaussian sum approximations: $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h) = \sum_{m=1}^{N_h} \alpha_m \mathcal{N}(\mathbf{h} \mid \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m)$ and $p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) = \sum_{n=1}^{N_R} \beta_n \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n)$. Then we have that

$$\begin{aligned} & \psi_h(\mathbf{h}, \mathbf{f}^R) p(\mathbf{h} \mid \mathbf{z}_{1:t}^h) p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) \\ &= \psi_h(\mathbf{h}, \mathbf{f}^R) \sum_{m=1}^{N_h} \alpha_m \mathcal{N}(\mathbf{h} \mid \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m) \sum_{n=1}^{N_R} \beta_n \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n), \end{aligned} \quad (\text{VI.1})$$

and as $N_h, N_R \rightarrow \infty$, we recover the densities $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ and $p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$ —see [1]. We note that the approximation $\sum_{n=1}^{N_R} \beta_n \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n)$ consists of only safe modes, since $\psi_f(\mathbf{f}^R, \mathbf{f})$ will assign near zero β to any unsafe modes. This is not true for the operator Gaussian mixture, since the operator may make choices that place the platform on a collision course.

Theorem 6.1 (LTB suboptimal): If either $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ or $p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$ is multimodal, then u_{LTB}^s is suboptimal with respect to agreeability, safety, and efficiency.

Proof: Let $\bar{\mathbf{h}} = \arg \max_{\mathbf{h}} p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ and $\bar{\mathbf{f}}^R = \arg \max_{\mathbf{f}^R, \mathbf{f}} p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$. Then we write Equation VI.1 as

$$\begin{aligned} & \psi_h(\mathbf{h}, \mathbf{f}^R) \sum_{m=1}^{N_h} \alpha_m \mathcal{N}(\mathbf{h} \mid \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m) \sum_{n=1}^{N_R} \beta_n \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n) \\ &= \psi_h(\mathbf{h}, \mathbf{f}^R) \left(\alpha^* \delta(\mathbf{h} - \bar{\mathbf{h}}) \beta^* \mathcal{N}(\mathbf{f}^R \mid \bar{\mathbf{f}}^R, \boldsymbol{\Sigma}^*) + \right. \\ & \quad \left. \sum_{m \neq *}^{N_h} \alpha_m \mathcal{N}(\mathbf{h} \mid \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m) \sum_{n \neq *}^{N_R} \beta_n \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n) \right) \end{aligned}$$

where we extracted the largest mixture components, and rewrote the summation using the notation $m \neq *$ to indicate that all the modes except the largest are present. If $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ and $p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$ are unimodal, then we recover linear blending as the optimally agreeable, safe, and efficient solution (Theorem 4.1). In general, however, the $\arg \max_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}}$ does not correspond to the linear blend.

We point out that this result is *independent of how $\psi(\mathbf{h}, \mathbf{f}^R)$ is chosen*. In other words, for any definition of “agreeability”, the suboptimality of linear blending still holds. ■

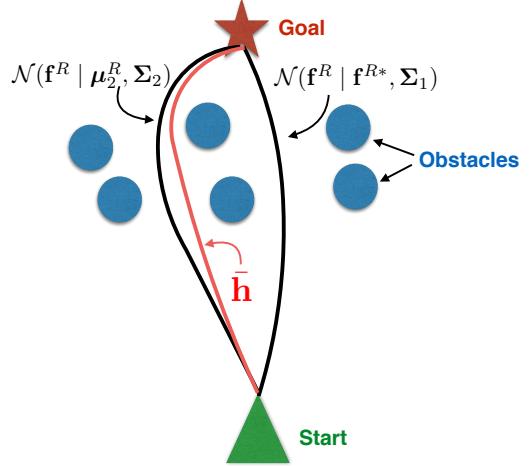


Fig. 3: One global autonomy optima at \mathbf{f}^{R*} and a safe but suboptimal autonomy mode at $\boldsymbol{\mu}_2^R$ through some obstacle field (additional autonomous modes exist but we leave them off for clarity). The operator’s unimodal predicted trajectory at $\bar{\mathbf{h}}$ is safe. Covariance functions removed for clarity.

To understand why this is important, first consider the illustration in Figure 3. Suppose that there are two autonomous safe modes through the obstacle field: $\mathcal{N}(\mathbf{f}^R \mid \mathbf{f}^{R*}, \boldsymbol{\Sigma}_1)$ and $\mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}_2^R, \boldsymbol{\Sigma}_2)$. Thus

$$\begin{aligned} & \psi_h(\mathbf{h}, \mathbf{f}^R) p(\mathbf{h} \mid \mathbf{z}_{1:t}^h) p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) \\ &= \psi_h(\bar{\mathbf{h}}, \mathbf{f}^R) [\beta_1 \mathcal{N}(\mathbf{f}^R \mid \mathbf{f}^{R*}, \boldsymbol{\Sigma}_1) + \beta_2 \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}_2^R, \boldsymbol{\Sigma}_2)] \\ &= \frac{\beta_1}{Z_1} \mathcal{N}(\mathbf{f}^R \mid \bar{\mathbf{f}}^{R*}, \bar{\boldsymbol{\Sigma}}_1) + \frac{\beta_2}{Z_2} \mathcal{N}(\mathbf{f}^R \mid \bar{\boldsymbol{\mu}}_2^R, \bar{\boldsymbol{\Sigma}}_2) \end{aligned}$$

where $\beta_1 > \beta_2$ (the first mode is the global optima). However

$$\begin{aligned} \frac{1}{Z_1} & \propto \exp \left(-\frac{1}{2} (\bar{\mathbf{h}} - \mathbf{f}^{R*})^\top (\gamma + \boldsymbol{\Sigma}_1)^{-1} (\bar{\mathbf{h}} - \mathbf{f}^{R*}) \right) \\ \frac{1}{Z_2} & \propto \exp \left(-\frac{1}{2} (\bar{\mathbf{h}} - \boldsymbol{\mu}_2^R)^\top (\gamma + \boldsymbol{\Sigma}_2)^{-1} (\bar{\mathbf{h}} - \boldsymbol{\mu}_2^R) \right), \end{aligned}$$

and so $1/Z_1$ is exponentially smaller than $1/Z_2$ since $\bar{\mathbf{h}} - \mathbf{f}^{R*}$ is much larger than $\bar{\mathbf{h}} - \boldsymbol{\mu}_2^R$. Thus, the probabilistic shared control in this situation is very close to both $\boldsymbol{\mu}_2^R$ and $\bar{\mathbf{h}}$.

Conversely, $u_{LTB}^s = K_h \bar{\mathbf{h}} + K_R \mathbf{f}^{R*}$. Here, if K_h is close to K_R , then u_{LTB}^s is unsafe. If $K_R \gg K_h$, then the autonomy overrides the operator’s (safe) choice. If $K_h \gg K_R$, then the operator is controlling the platform, and is thus not being assisted. This begs the following question: what heuristic should be employed to choose K_h and K_R ? By carrying multiple modes (as in probabilistic shared control), we bypass this dilemma, since heuristics are never invoked. Basic rules of probability theory (namely, the normalizing factor) determine the best choice. In other words, probabilistic shared control is able to determine the shared control in a data driven way rather than through anecdote.

In Figure 4 the operator chooses an *unsafe* trajectory. For linear trajectory blending, $u_{LTB}^s = K_h \bar{\mathbf{h}} + K_R \mathbf{f}^{R*}$, and so we

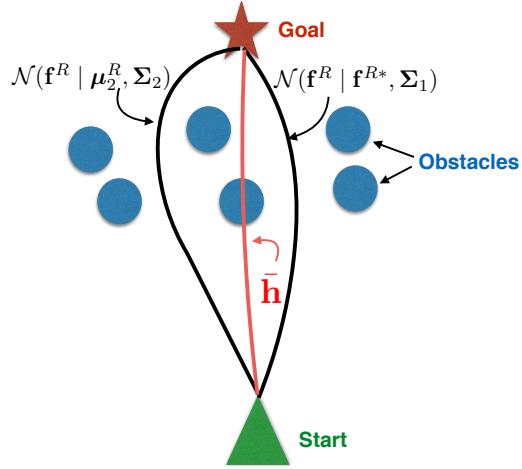


Fig. 4: One global autonomy optima at f^{R*} and a safe but suboptimal autonomy mode at μ_2^R through some obstacle field (additional autonomous modes exist but we leave them off for clarity). The operator's unimodal predicted trajectory at \bar{h} is *unsafe*. Covariance functions removed for clarity.

must choose our heuristics such that $K_R \gg K_h$ in order to avoid collision—that is, we must insert logic that overrides the operator when the operator makes unsafe decisions.

In contrast, consider

$$\begin{aligned} \psi_h(\mathbf{h}, \mathbf{f}^R) p(\mathbf{h} | \mathbf{z}_{1:t}^h) p(\mathbf{f}^R, \mathbf{f} | \bar{\mathbf{z}}_{1:t}) &= \psi_h(\mathbf{h}, \mathbf{f}^R) p(\mathbf{h} | \mathbf{z}_{1:t}^h) \times \\ [\beta_1 \mathcal{N}(\mathbf{f}^R | \mathbf{f}^{R*}, \Sigma_1) + \beta_2 \mathcal{N}(\mathbf{f}^R | \mu_2^R, \Sigma_2)] \\ &= \psi_h(\mathbf{h}, \mathbf{f}^R) \mathcal{N}(\mathbf{h} | \mu_h, \Sigma_h) \times \\ [\beta_1 \mathcal{N}(\mathbf{f}^R | \mathbf{f}^{R*}, \Sigma_1) + \beta_2 \mathcal{N}(\mathbf{f}^R | \mu_2^R, \Sigma_2)] \end{aligned}$$

where we have maintained both autonomous modes and a unimodal distribution over \mathbf{h} . In this situation β_1 is close to β_2 , and the difference between the operator mean and either of the autonomy means are nearly the same, so Z_1 is close to Z_2 . However, both Σ_1 and Σ_2 are both fairly narrow (otherwise, they are not safe modes), and because there is flexibility in $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$, the MAP value is close to either \mathbf{f}^{R*} or μ_2^R . Because the operator is treated probabilistically, heuristics are not used to detect poor operator choices: the autonomy assists the operator by blending near the tails of $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$.

In Figure 5, the operator generates ambiguous data about how he wishes to move between the start and the goal, and so $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ is bimodal. Linear blending produces $u_{LTB}^s = K_h \bar{h} + K_R \mathbf{f}^{R*}$ and so we end up with a situation very similar to that discussed in Figure 3—an autonomy and an operator that are needlessly in disagreement, and thus difficult to disambiguate with K_h and K_R . To be fair, if \bar{h} and \mathbf{f}^{R*} happen to lie close to one another, then linear blending provides the optimal solution. But for multimodal autonomous and operator distributions, such a situation is the exception rather than the rule (this exception requiring that the operator make globally optimal decisions).

For probabilistic shared control, the shared control is likely a trajectory near \bar{h} and μ_2^R —a solution that is both safe and respects the operator's desires (although, depending on the weights α_m and β_n the solution may be near μ_2^h and μ_3^R ,

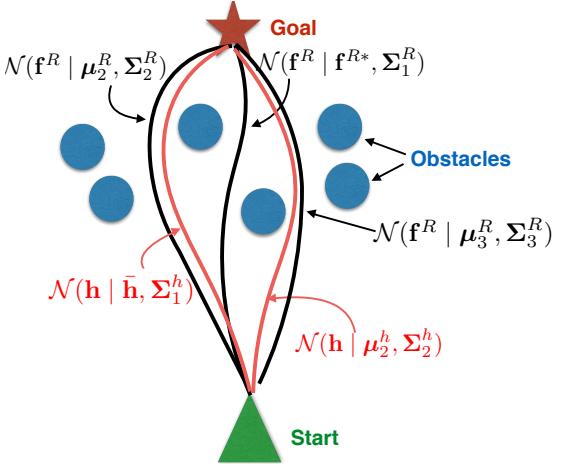


Fig. 5: One global autonomy optima at f^{R*} and two safe but suboptimal autonomy modes at μ_2^R and μ_3^R through some obstacle field. The operator's bimodal predicted trajectory centered at \bar{h} and μ_2^h . Covariance functions removed for clarity.

a solution that reflects the operator's desires and is still safe).

These figures illustrate Theorem 6.1: that linear blending is suboptimal when either the operator or the autonomy model is multimodal; anecdotally, this is a result of restricting the blend to a single autonomous decision and a single operator trajectory, which can be in conflict when nontrivial modes are not reasoned over.

VII. CONCLUSIONS

We presented a probabilistic formalism for shared control, and showed how the state of the art in shared control, linear blending, and a trajectory based generalization, are special cases of our probabilistic approach. Further, we showed that linear blending is prone to statistical inconsistencies, but that the probabilistic approach is not. Finally, we proved that linear blending is suboptimal with respect to agreeability, safety and efficiency. Importantly, no experiments were performed, and so questions remain as to the viability of this approach.

REFERENCES

- [1] D. Alspach and H. Sorenson. Nonlinear bayesian estimation using gaussian sum approximations. *IEEE Trans. on Automatic Control*, 1972.
- [2] A. Bejczy. Recent developments in advanced teleoperation at jpl. In *IEEE International Conference on Advanced Robotics*, 1991.
- [3] C. Bishop. *Pattern Recognition and Machine Learning*. Springer Science+Business Media, LLC, New York, NY, 2006.
- [4] J. Burke and R. Murphy. From remote tool to shared roles. In *IEEE Robotics and Automation Magazine*, 2008.
- [5] T. Carlson and Y. Demiris. Collaborative control for a robotic wheelchair: Collaborative control for a robotic wheelchair: Evaluation of performance, attention and workload. *IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS—PART B: CYBERNETICS*, 2012.
- [6] E. Demeester, A. Huntemann, E. V. Poorten, and J. D. Schutter. MI, map and greedy pomdp shared control: comparison of wheelchair navigation assistance for switch interfaces. In *ISER*, 2012.
- [7] E. Demeester, A. Huntemann, D. Vanhooydonck, F. Vanacker, H. V. Brussel, and M. Nuttin. User-adapted plan recognition and user-adapted shared control: A bayesian approach to semi-autonomous wheelchair driving. In *Autonomous Robots*, 2008.
- [8] A. Dragan and S. Srinivasa. Formalizing assistive teleoperation. *Robotics: Science and Systems*, 2012.

- [9] A. Dragan and S. Srinivasa. A policy blending formalism for shared control. *International Journal of Robotics Research*, 2013.
- [10] D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics and Automation Magazine*, 1997.
- [11] J. Fu and U. Topcu. Probably approximately correct mdp learning and control with temporal logic constraints. In *RSS*, 2014.
- [12] S. Gulati, C. Jhurani, B. Kuipers, and R. Longoria. A framework for planning comfortable and customizable motion of an assistive mobile robot. In *IROS*, 2009.
- [13] B. Hardin and M. Goodrich. On using mixed-initiative control: A perspective for managing large-scale robotic teams. In *ACM/IEEE International Conference on Human-Robot Interaction*, 2009.
- [14] A. Huntemann, E. Demeester, and E. V. Poorten. Probabilistic approach to recognize local navigation plans by fusing past driving information with a personalized user model. In *ICRA*, 2013.
- [15] P. Inigo-Blasco, F. D. del Rio, S. V. Diaz, and D. C. Muniz. The shared control dynamic window approach for non-holonomic semi-autonomous robots. In *International Symposium on Robotics/Robotik*, 2014.
- [16] W. B. Knox and P. Stone. Tamer: Training an agent manually via evaluative reinforcement. In *ICDL*, 2008.
- [17] A. Lopes, U. Nunes, and L. Vaz. Assisted navigation based on shared-control, using discrete and sparse human-machine interfaces. In *IEEE Engineering in Medicine and Biology Society*, 2010.
- [18] A. M. Morison, D. D. Woods, and T. Murphy. Human-robot interaction as extending human perception to new scales. In *Handbook of Applied Perception*. Cambridge University Press, in press.
- [19] G. Peinado, C. Urdiales, J. Peula, M. Fdez-Carmona, R. Annicchiarico, F. Sandoval, and C. Caltagirone. Navigation skills based profiling for collaborative wheelchair control. In *ICRA*, 2011.
- [20] A. Poncea, C. Urdiales, E. Perez, and F. Sandoval. A new efficiency-weighted strategy for continuous human/robot cooperation in navigation. *IEEE Trans. on SMC—Part A: Systems and Humans*, 2009.
- [21] P. Trautman and A. Krause. Unfreezing the robot: Navigation in dense interacting crowds. In *IROS*, 2010.
- [22] P. Trautman, J. Ma, R. Murray, and A. Krause. Robot navigation in dense crowds: the case for cooperation. In *ICRA*, 2013.
- [23] P. Trautman, J. Ma, R. M. Murray, and A. Krause. Robot navigation in dense human crowds: Statistical models and experimental studies of human robot cooperation. *IJRR*, 2015.
- [24] I. Ulrich and J. Borenstein. Vfh+: Reliable obstacle avoidance for fast mobile robots. In *ICRA*, 1998.
- [25] C. Urdiales, E. Pérez, G. Peinado, M. Fdez-Carmona, J. Peula, R. Annicchiarico, F. Sandoval, and C. Caltagirone. On the construction of a skill-based wheelchair navigation profile. *IEEE TRANSACTIONS ON NEURAL SYSTEMS AND REHABILITATION ENGINEERING*, 2013.
- [26] C. Urdiales, J. Peula, M. Fdez-Carmona, C. Barrué, E. Pérez, I. Sánchez-Tato, J. del Toro, F. Galluppi, U. Cortés, and R. Annicchiarico. A new multi-criteria optimization strategy for shared control in wheelchair assisted navigation. *Autonomous Robots*, 2011.
- [27] H. Wang and X. Liu. Adaptive shared control for a novel mobile assistive robot. *IEEE/ASME TRANSACTIONS ON MECHATRONICS*, 2014.
- [28] M. Wang and J. Liu. Interactive control for internet-based mobile robot teleoperation. *Robotics and Autonomous Systems*, 2005.
- [29] P. Welinder, S. Branson, S. Belongie, and P. Perona. The multidimensional wisdom of crowds. In *NIPS*, 2010.
- [30] H. Yu, M. Spenko, and S. Dubowsky. An adaptive shared control system for an intelligent mobility aid for the elderly. *Autonomous Robots*, 2003.